REPORT DOCUMENTATION PAGE	*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
27-12-2011	Final	01-01-09 to 30-SEP-11

4. TITLE AND SUBTITLE

Large-Scale Online and Real-Time Optimization Problems under Uncertainty

5a. CONTRACT NUMBER

5b. GRANT NUMBER

N00014-09-1-0326

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Jaillet, Patrick

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Massachusetts Institute of Technology
Office of Sponsored Programs
77 Massachusetts Ave
Cambridge, MA 02139-4301

8. PERFORMING ORGANIZATION REPORT NUMBER

6918983

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
875 North Randolph Street
Arlington, VA 22203-1995

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution A. Approved for pubic release: distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER *(Include area code)*

LARGE-SCALE ONLINE AND REAL-TIME OPTIMIZATION PROBLEMS UNDER UNCERTAINTY

Patrick Jaillet, MIT

ONR grant N00014-09-1-0326, 1/1/09 - 12/31/11
Federal Identifier: 311 - Dr. Donald Wagner
Final report: December 2011

1 Overview of our long-range research objectives

The main focus of our research is on the fundamental aspects of decision making in the context of uncertain (and possibly very large) data sets revealed in an online fashion. More specifically, we are interested in the intersection and interplay of three main phenomena (incomplete and uncertain data, online decisions with or without real-time restrictions, and extremely large data sets), and the corresponding fundamental questions when facing a problem exhibiting one are several of these phenomena:

- How to properly model and quantify the degree of uncertainty in such a problem, and its impact on our ability to "solve" it?
- How to quantify the value of additional (deterministic and/or probabilistic) information for solving such problems? Is it worth (and how much) to seek additional information?
- How to design provably good algorithms for such problems?

There are many motivating examples for such problems. In the proposal that led to this funded research, we listed examples in routing and scheduling as well as applications in which the basic decisions were to assign goods, items, or requests to clients, bins or servers (assignment or allocation problems). In a recent internal ONR white paper ([1]) we have highlighted applications associated with the deployment of autonomous multi-agent systems for spatial exploration and information harvesting.

The major objectives of this long-ranging research involve the definition and rigorous mathematical analysis of canonical models which would capture the essence of this new class of problems. The overall purpose is to systematically (i) analyze and integrate the most promising paradigms/solution strategies for dealing with these problems, (ii) formulate adequate corresponding models, and (iii) develop, analyze, and implement algorithms for solving them.

2 Results from grant N00014-09-1-0326

2.1 Research emphasis

In our research, we have concentrated such efforts on two basic canonical online combinatorial optimization problems, one based on various online versions of the traveling salesman

problem (TSP) (and its generalization to multi-servers), and one based on online versions of the bipartite matching problem (and its generalization to various weighted cases).

2.2 Summary of results

We have made significant progress and ended up on target for what we had proposed to do. Below we provide technical details about what we have done so far as contained in various written papers. We then list some additional information. In a final section we highlight the main remaining open research questions as well as the new ones which came of this research, and which we have proposed to expand upon as part of the research described in a renewal grant whose funding was recently approved (grant N00014-12-1-0033).

2.3 Technical details of results

1. Jaillet, P. and X. Lu, "Online Traveling Salesman Problems with Service Flexibility", *Networks*, 58, 137-146, 2011. ([5])

In this paper, we assume that there is a penalty for not serving a request. Requests for visit of points in the metric space are revealed over time to a server, initially at a given origin, who must decide in an online fashion which requests to serve in order to minimize the time to serve all accepted requests plus the sum of the penalties associated with the rejected requests. We first look at the special case of the non-negative real line. After providing a polynomial time algorithm for the offline version of the problem, we propose and prove the optimality of a 2-competitive polynomial time online algorithm based on re-optimization approaches. We also consider the impact of advanced information (lookahead) on this optimal competitive ratio. We then consider the generalizations of these results to the case of the real line. We show that the previous algorithm can be extended to an optimal 2-competitive online algorithm. Finally we consider the case of a general metric space and propose an original c-competitive online algorithm, where $c = \frac{\sqrt{17}+5}{4} \approx 2.28$. We also give a polynomial-time $(1.5\rho + 1)$-competitive online algorithm which uses a polynomial-time ρ-approximation for the offline problem.

2. Jaillet, P. and X. Lu, "Online Traveling Salesman Problems and Hamiltonian Path Problems with Rejection Options," Working Paper, MIT, initial version 2010, expanded version for submission to *Mathematics of Operations Research*, 2011. ([4])

In this paper, we consider online versions of the TSP and Hamiltonian Path Problem for which requests continue to be revealed over time to a server, initially at a given origin, who must decide which requests to serve in order to minimize the time to serve all accepted requests plus the sum of the penalties associated with the rejected requests. In the first online version of this problem, hereafter called the basic version, we assume that the server's decision to accept or reject a request can be made any time after its release date. In the second online version of this problem, hereafter called the real-time version, we assume that the server's decision to accept or reject a request must be made exactly at its release date. In this paper, we first provide an optimal 2-competitive online algorithm for the basic version of the TSP in a general metric space, improving the 7/3-competitive online algorithm of Ausiello et al.[Inf.Proc.Letters, 2008], and the 2.28-competitive online algorithm described

in our first paper [5]. For the HPP, we give a 2-competitive online algorithm on the non-negative real line, and a $(\sqrt{2}+1)$-competitive online algorithm on general metric spaces.

We then consider the real-time version of the TSP. We provide a best possible 2.5-competitive polynomial time online algorithm on the non-negative real line. On the real line, we prove a lower bound of 2.64 on any competitive ratios and give a 3-competitive online algorithm. We then consider the case of a general metric space and prove a $\Omega(\sqrt{\ln n})$ lower bound on any competitive ratios, and describe an asymptotically best possible $O(\sqrt{\ln n})$-competitive online algorithm.

3. Therkelsen, C.W. "TSP with with Flexible Service and Multiple Servers," Bachelor thesis, Electrical Engineering and Computer Science, MIT, May 2011. Supervised by P. Jaillet. ([8])

This work is concerned with the Online Traveling Salesman Problem (TSP) with Flexible Service and with Multiple Servers. We consider two algorithms; one of which is 2-competitive and best-possible, and one of which is 2.5-competitive. We also consider the asymptotic performance of the 2.5-competitive algorithm. We show that the 2.5-competitive algorithm is asymptotically optimal (i.e. the ratio with the value of the 2.5-competitive algorithm in the numerator and the value of the optimal offline solution in the denominator converges to 1 almost surely). Finally we consider the performance of these two algorithms empirically through simulations on medium size instances of the problem given certain stochastic conditions. The simulations indicate that the 2.5-competitive algorithm performs better than the 2-competitive and best-possible algorithm in practice, from which we conclude that while the 2-competitive algorithm might be optimal if the problem instance is determined by a malevolent agent, other algorithms will perform better if the problem instance can be modeled as a random process.

4. Jaillet, P. and X. Lu, "Online Generalized Assignment Problems," Working Paper, MIT, initial version 2010, new expanded version for submission to *Information Processing Letter*, 2011. ([2])

Concentrating on a series of applications where dynamic resource allocations arise naturally, we formulate an online version of a generalized assignment problem, where objects are to be assigned to buyers with some budget constraints. We propose, analyze, and empirically compare two online algorithms, one based on a greedy principle and the other one on a primal-dual approach. In particular, we show that for a large class of instances, both algorithms have an asymptotically optimal competitive ratio of 1/2, and prove that they are also optimal among randomized algorithms. We also provide a general technique in order to transform a c-competitive algorithm on this class of instances to a c'-competitive algorithm on a more general class of instances, with proven tight bounds on c', depending on c and other parameters defining the new class of instances.

5. Jaillet, P. and X. Lu, "Online Stochastic Matching: New Algorithms with Better Bounds," Working Paper, expanded version for submission to *Mathematics of Operations Research*, 2011. ([3])

In this most recent series of results, we consider the online bipartite matching problem. In this setting, one set of the bipartite graph is fixed and known in advance (the servers). The other set is composed of various request types (a type describes which subset of servers can serve the request). Requests come one at a time and types are discovered upon arrivals.

The task is to assign each request to a capable server (or to discard it) immediately upon its arrival. In the adversarial model, the ranking algorithm of Karp et al. [STOC, 1990] provides a best possible randomized algorithm with competitive ratio $1 - 1/e \approx 0.632$ for this problem. In the random permutation model, when the set of requests is arbitrary, but the order of the sequence is random, Goel and Mehta [SODA, 2008] show that a greedy algorithm matches this bound. In the stochastic i.i.d. model, when requests are drawn repeatedly and independently from a known probability distribution over the different request types, one can do better (Feldman et al. [FOCS, 2009 provide a 0.670-competitive algorithm, Manshadi et al. [SODA 2011] propose a 0.702-competitive algorithm). In this paper we consider a novel class of online algorithms for the i.i.d. model which improve on all these bounds and which use computationally efficient offline procedures. We get a 0.729-competitive. Compared to the results of Manshadi et al. our algorithms don't rely on the need to first run a (possibly computationally intensive) Monte Carlo sampling procedure, but rather are based on the solution of simple linear programs which have the basic simple structure of maximum flow problems. We also show that our techniques can be applied to other related problems such as the online stochastic vertex-weighted bipartite matching problem as defined in Aggarwal et al. [SODA 2011], for which we obtain a 0.725-competitive algorithm under the general stochastic i.i.d. model (again showing one can do better than the best possible bound of $1 - 1/e$ in the adversarial model, as established Aggarwal et al).

6. Recent papers of our post-doc Rico Zenklusen have also been partially supported by this grant, and led to the proposed consideration in our renewal grant of a new class of canonical online models based on various generalizations of the classical secretary problem (with matroidal constraints): [6]; [7]; and [9].

2.4 Funded personnel and facilities

Personnel: The personnel funded by this grant has included:

> Prof. Jaillet, PI: summer month support in 2010 and 2011.
>
> Xin Lu, ORC doctoral student: RA stipend + tuition for Fall 09, Spring 10, Summer 10, Fall 10, Spring 11, Summer 11 and part of Fall 11. Tasks: Theoretical development behind most results funded from this research; main part of his doctoral dissertation.
>
> Maokai Lin, ORC doctoral student: RA stipend + tuition for Fall 09. Tasks: Development of a simulation framework for the testing of online algorithms for the online TSP.
>
> Rico Zenklusen, post-doc, Math and EECS: part-time salary Summer and Fall 2011. Tasks: Collaborate on the theoretical development and expansion of the research tools.

Other personnel working on this research without direct funding included:

> Christian Therkelsen, EECS undergraduate senior; Spring 2011; thesis supervised by Prof Jaillet.

Facilities: The main facilities available for conducting this research included offices and computing capabilities within the Laboratory for Information and Decision Systems (LIDS) and the Operations Research Center (ORC) at MIT.

2.5 Presentations and other information

Presentations: Our research has been presented at major national and international conferences (such as INFORMS annual meetings, TRISTAN international meetings) as well as in select specialized workshops (Schloss Dagstuhl, Germany; Norwegian Research Council, Molde, Norway; AFOSR-DDDAS, Washington; ONR-IPAM, UCLA). The research has also been presented at invited seminars at various universities.

Honors and recognitions: Over the duration of this grant, the PI received the Dugald C. Jackson Professorship in the Department of Electrical Engineering and Computer Science, MIT; became Co-Director of the MIT Operations Research Center in January 2010; and is currently serving as President of the Transportation Science and Logistics (TSL) Society of INFORMS.

Derived outcomes: From conference presentations and workshop presentations, as described above, contacts have been made which have resulted in the following follow-up activities:

Co-supervision of an MIT SM ORC thesis of a Draper Fellow from the Navy, Jacob Cates, entitled "Route Optimization Under Uncertainty for Unmanned Underwater Vehicles".

Collaboration with a team from Lockheed Martin Advanced Technology Laboratories in the writing of a white paper and forthcoming participation as a consultant in a recently funded proposal from Lockheed Martin ATL in response to ONR SN 11-SN-0012 (PI Ray Yuan: Probabilistic Auction for Distributed, Robust, Resource, Optimization).

3 Conclusion

In our research we have looked at some of the fundamental intellectual challenges behind (1) the integration of (i) incomplete and uncertain data, (ii) short time restriction for decisions, and (iii) extremely large data sets; (2) the development of appropriate canonical models for considering these multi-facet issues, and (3) the design of new algorithmic solutions and paradigms for addressing these canonical models. More specifically we have introduced and analyzed successfully along these lines online versions of two basic combinatorial optimization problems, the traveling salesman problem (TSP) and the bipartite bipartite matching.

Main open questions and issues in this ambitious agenda have yet to be fully addressed and understood. Two fundamental ones include (a) the impact of limited access to past information; and (b) the proper inclusion of stochastic information in a robust way. Some new important research questions have surfaced during this past investigation (for example the modeling of learning mechanism for online optimization). In a renewal (and recently) funded proposal (N00014-12-1-0033), we propose to address all these research issues for our two basic canonical models and their variations, as well as for a new class of canonical models based on various generalizations of the classical secretary problem.

References

[1] P. Jaillet. Data-driven and online optimization under uncertainty: Opportunities, challenges, and research needs, 2011. White Paper, ONR internal.

[2] P. Jaillet and X. Lu. Online generalized assignment problems. working paper, Massachusetts Institute of Technology, 2011.

[3] P. Jaillet and X. Lu. Online stochastic matching: New algorithms with better bounds. working paper, Massachusetts Institute of Technology, 2011.

[4] P. Jaillet and X. Lu. Online traveling salesman problems and hamiltonian path problems with immediate decisions. working paper, Massachusetts Institute of Technology, 2011.

[5] P. Jaillet and X. Lu. Online traveling salesman problems with service flexibility. *Networks*, 58:137–146, 2011.

[6] K. Schüpbach and R. Zenklusen. An adaptive routing approach for personal rapid transit, 2011. Submitted.

[7] K. Schüpbach and R. Zenklusen. Approximation algorithms for conflict-free vehicle routing. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, pages 640–651, 2011.

[8] C.W Therkelsen. Tsp with with flexible service and multiple servers, 2011. Undergraduate thesis, EECS, MIT.

[9] R. Zenklusen. Matroidal degree-bounded minimum spanning trees. In *Proceedings of the 23rd Annual ACM -SIAM Symposium on Discrete Algorithms (SODA)*, 2012. To appear.